



Europäisches Patentamt
European Patent Office
Office européen des brevets



⑪ Publication number:

0 492 106 A1

⑫

EUROPEAN PATENT APPLICATION

⑬ Application number: 91119207.8

⑮ Int. Cl. 5: G01C 22/02, G11C 16/06,
G06F 11/20

⑭ Date of filing: 12.11.91

⑯ Priority: 28.12.90 US 635910

⑰ Inventor: Pricer, Wilbur Davis

⑰ Date of publication of application:
01.07.92 Bulletin 92/27

58 Van Patten Parkway
Burlington, Vermont 05452(US)

⑱ Designated Contracting States:
DE FR GB

⑲ Representative: Herzog, Friedrich Joachim,
Dipl.-Ing.

⑲ Applicant: INTERNATIONAL BUSINESS
MACHINES CORPORATION

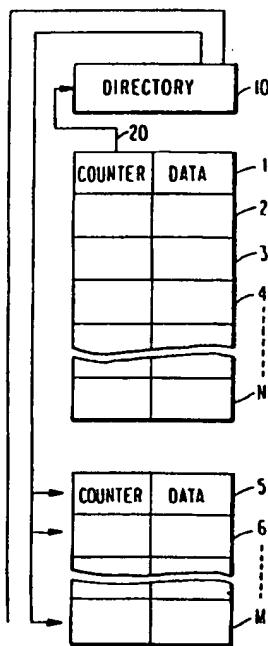
IBM Deutschland GmbH, Patentwesen und
Urheberrecht, Schönaicher Strasse 220
W-7030 Böblingen(DE)

Armonk, NY 10504(US)

⑳ Endurance management for solid state files.

⑳ A nonvolatile storage system including several data storage arrays; a counter storage array for storing values indicating the remaining useful life of each of the data storage arrays; and a device for controlling access to each of the data storage arrays as a function of the values stored in each counter array. The counter array has an appreciably longer useful life than any of the data storage arrays.

FIG. 1



EP 0 492 106 A1

The present invention relates to an apparatus and method for counting the number of erase/write cycles that have occurred in nonvolatile files, and for reallocating a file when the counted number of erase/write cycles of the file approaches its endurance limit.

The cost per bit of magnetic memory, such as magnetic disks, is decreasing. However, the cost per bit of nonvolatile memory, such as EEPROMs, is decreasing at an even faster rate. Presently, most systems do not employ nonvolatile devices due to their limited endurance, which in the best EEPROM technologies is at most approximately 10^7 erase/write cycles. This is due to the breakdown of tunneling dielectrics. When nonvolatile devices are employed, severe limitations are placed on the circuit design. Specifically, such a system must be configured to limit the number of erase/write cycles experienced by the nonvolatile devices. The following are examples of such systems.

U.S. Patent 4,803,707 to Cordan Jr. discloses a vehicle odometer system for counting pulses representing vehicle mileage, and for storing the signal count during periods when the vehicle power is disabled. In this system, a group of nonvolatile memory cells is used to store binary values for odometer readings up to 10,000 miles. For odometer readings greater than 10,000, a new set of memory cells is accessed. Thus, each group of memory cells is cycled up to only 10,000 times which is below the endurance limit. The higher order bits of each memory block, which are not cycled as often as the lower-order bits, serve as a pointer for the new set of memory cells.

In order to ensure the integrity of the stored data, this system requires that each group of nonvolatile memory cells be reallocated well below the maximum endurance limit of the cells. However, such a reallocation scheme is not practical in a system containing a large number memory blocks because the memory blocks would be reallocated too often.

U.S. Patent No. 4,528,683 to Henry discloses an odometer system which employs a highly particularised five-bit word per decimal position counting scheme, along with a multi-level multiplexer to move the decimal ones, tens, hundreds and thousands position through a nonvolatile memory array. This system allows erase/cycle "wear" effects of the cells to be spread throughout the memory array, thereby preventing any one group of cells from wearing out too quickly.

U.S. 4,663,770 to Murray et al discloses a system including nonvolatile counter devices. In order to distribute the erase/write cycle wear effects among the various counter devices, counter usage is successively shifted. While the scheme of this system allows for distribution of the erase/write wear effects among the nonvolatile devices within the system, the coding and decoding for this distribution requires a complex hardware configuration.

It is an object of the invention to provide a nonvolatile memory system containing a large number of primary large memory blocks, and a minimum number of identical "substitutional" or reallocation memory blocks.

It is another object of the invention to provide a system which utilizes nonvolatile devices and which is relatively simple in design.

It is still another object of the invention to provide a system which ensures the integrity of data stored in nonvolatile memory even though the memory is subjected to a large number of erase/write cycles.

The present invention as laid down in the claims accomplishes these and other objects in advantageous manner. The invention provides this by counting the number of erase/write cycles that a nonvolatile data block experiences, and reallocating the data block when the counted number of erase/write cycles approaches a predetermined number representing the endurance limit of the nonvolatile data block.

In the following the invention will be described in more detail in connection with the embodiments shown in the drawing, in which

- 45 Fig. 1 is a block diagram showing a memory system according to the invention;
- Fig. 2 is a block diagram of a counter utilized in the system shown in Fig. 1;
- Fig. 3 is a circuit block diagram for illustrating a counting operation according to the invention, and
- Fig. 4 is a flowchart showing the operation according to the invention.

50 PREFERRED EMBODIMENTS OF THE INVENTION

As shown in Fig. 1, the memory system according to the invention includes several primary counter/data registers 1, 2, 3, 4,...N, several substitutional counter/data registers 5, 6,...M; and a directory processor 10.

55 Each of the primary and substitutional counter/data registers includes a large nonvolatile memory block, e.g., which contains 2×10^3 bytes, and a counter for counting each erase/write cycle that the associated memory block experiences. During an erase/write cycle for a particular memory block, the counter associated with that particular memory block is first read. If the read count is below a predetermined

number, such as 1.0×10^7 which is the erase/write cycle endurance limit for a nonvolatile memory block, then the counter is incremented by one. However, if the read count is above the erase/write cycle endurance limit, then the counter outputs a reallocation signal, via signal line 20, to the directory processor 10 indicating that the memory block should be reallocated. Upon receiving the reallocation signal from a 5 counter, the directory processor 10 downloads the data in the memory block, whose erase/write cycle endurance limit has been reached, into one of the substitutional counter/data registers 5, 6,...M which has not yet been accessed, i.e., subjected to any erase/write cycles. Since the substitutional counter/data register has not yet been previously accessed, the reallocated data can now be safely subjected to another 1.0×10^7 erase/write cycles.

10 The directory processor 10 functions to keep track of which primary and substitutional counter/data registers contain data. For example, if primary counter/data register 1 sends a reallocation signal to the directory processor 10, the directory processor 10 will determine which of the substitutional counter/data registers 5,6...M do not contain data, i.e., have not been accessed. If, for example, substitutional counter/data register 6 has not been accessed, then the directory processor 10 will reallocate the data from 15 primary counter/data register 1 to the substitutional counter/data register 6. The address of counter/data register 6 will now serve as the address for the data previously stored in primary counter/data register 1.

15 Fig. 2 shows the counter of the counter/data registers in greater detail. As shown in Fig. 2, the counter includes a twenty- eight bit binary register 30 which stores the high-order bits of the counted data, a pair of nonvolatile linear registers 40 and 50 which store the lower-order bits of the counted data, and a comparator 20 60 for comparing the data stored in the linear registers 40 and 50.

25 The binary register includes twenty-two bits for storing data, and six-bits for an ECC (Error Correction Code). Each of the linear registers stores twelve-bits of counted data. The following Table exemplifies the count sequence of the counter shown in Fig. 2 between erase/write cycles N and N + 4.

25

30

E/W Cycle	Binary Register	Linear Registers
N	10010	000111
N + 1	10010	001111
N + 2	10010	011111
N + 3	10010	111111
N + 4	10011	000000

35 In this example, the nonvolatile counter bits experience 1/12 as many erase/write cycles as do the data bits, and therefore the counter bits do not deteriorate as quickly as do the data bits of the associated memory block, thereby ensuring an accurate count of the number of erase/write cycles of the associated data block.

40 In Fig. 2 the comparator 60 compares the data contained in the linear registers 40 and 50. Since the linear registers 40, 50 are identical and count the same erase/write cycles, the counted data contained therein should always be the same. However, in the event of an error in either of the two linear registers, the higher count is assumed correct. In the event of an error in the binary register, the error is corrected by the ECC. This technique could be extended to higher orders of redundancy, if desired.

45 Fig. 3 shows the operation of counting the linear registers. The data bits are directly connected to a driver 70, and the linear counter bits are connected to the driver 70 via the data bits and a switch 80. During all accesses, the counter bits are read first. During the erase portion of the erase/write cycle, If the linear counters are not full, the linear counter bits will be disconnected from the driver 70 by operating switch 80 to the "UP" position, as shown in

50 Fig. 3. During the write portion of the cycle, only the next linear bit will be written. If the linear counter bits are full, then the counter bits will remain connected to the driver for the entire erase/write cycle by operating switch 80 to the "DOWN" position. During the erase/write cycle, the linear register will be reset to all zeros, and the binary counter counted-up by one.

55 Fig. 4 shows the operation of the memory system. As shown in Fig. 4, it is first determined in step S1 whether or not an erase/write cycle has been requested for a particular data block. If an erase/write cycle has been requested, then the counter associated with that particular memory block is read at step S2. The linear registers 40 and 50 are then compared at step S3 to determine if they are equal to each other. As indicated above, the linear registers 40 and 50 are identical, and should always contain the same data. If it is determined in step S3 that the linear registers 40 and 50 are equal to each other, then the process proceeds to step S4 where it is determined whether the counter is full, i.e., whether the counter has reached the predetermined endurance limit of the memory block. If it is determined that the counter is not full at

step S4, then the process proceeds to step S5 where it is determined whether or not any error has occurred by examining the ECC of the binary register. If it is determined that no errors have occurred, then the counter is incremented by 1, and the process proceeds back to step S1.

If in step S3 it is determined that the linear registers 40 and 50 are not equal to each other, then it is assumed that an error has occurred, and the process proceeds to step S7 where the higher data count of the two registers is taken as the correct count. The process then proceeds to step S8 where it is determined whether or not the counter is full. If the counter is not full, then the process proceeds to step S9 where the counter is incremented by a value X which is some significant fraction of the counter's full capacity, because it was determined in step S3 that an error occurred, i.e., the linear registers were not equal to each other in step S3. Such an error may mean that the nonvolatile memory cells are beginning to wear out. By incrementing the counter by the value X, the count is significantly increased. As a result, the memory block associated with the counter will be reallocated sooner.

If in steps S4 or S8 it is determined that the counter is full, then the data contained in the memory block is reallocated to a empty or spare substitutional data/counter memory block, as discussed above. If in step S5 it is determined from the ECC that an error has occurred, then the count is incremented by the value X. Again, since an error has occurred, this may mean that the nonvolatile memory cells are beginning to wear out. By incrementing the counter with the value X which is some significant fraction of the total count, the data contained in the associated data block will be reallocated sooner.

In another embodiment of the invention, the linear registers are replaced with a random number generator programmed to generate an output having a one-in-twelve chance of being a "1" or HIGH output in any erase/write cycle. Upon generation of a "1" output, the binary register would be counted-up. While in this embodiment the exact count of the write/erase cycles would no longer be absolutely correct, the count would be statistically accurate enough. Alternatively, the random number generator may be a pseudo-random number generator such a linear feedback shift register of long prime number length, properly weighted.

In another embodiment of the invention, the nonvolatile memory blocks are tested by subjecting the memory blocks to relatively short erase/write cycles to determine which memory blocks are likely to fail first, and which memory blocks have a better than average endurance to erase/write cycles. Those memory blocks which are likely to fail first could be loaded with a pre-count so that if these blocks are subjected to substantially more than average usage, the blocks would be reallocated early, i.e., earlier than if the counter was not loaded with a pre-count. On the other hand, those memory blocks determined to have a better than average endurance to erase/write cycles could be used as substitutional blocks. These high endurance memory blocks would later be substituted for high usage addresses. Thus, the strongest storage cells would subsequently be used for the highest usage addresses. Further, the counters of the system could be combined with a redundancy bad-block scheme wherein those memory blocks determined to have extremely poor endurance to erase/write cycles are preloaded with a full count. This could be done when the memory is first manufactured, or later on a periodic basis during use.

In another embodiment of the invention, the counters store count values corresponding to endurance-limiting phenomena other than numbers of erase/write cycles. For example, a sensor could be used to indicate that the memory is being stressed by the application of high power surges or voltages, or other temporary endurance-detracting environmental conditions such as high temperatures. A count could be generated that varies as a function of the amount of time the stress conditions are experienced, the absolute value of such stress conditions, and the count already stored in the counter (i.e., as the counter value increases, the above stress conditions will have a greater effect upon endurance). Another alternative is to generate counts that reflect other more permanent environmental endurance detractors, such as high humidity, low pressure, ambients having high particulate counts, etc. Another alternative is to monitor the validity of the data from the memory using error correction codes (ECC) to increment the counter for memory blocks experiencing soft errors. Again, the count could be weighed such that the count is higher as the number of soft errors increases, or as the soft errors occur later in time. Yet another alternative is to identify "key" memory blocks that will be used most often and/or will store critical information. The counters associated with these blocks could be incremented prior to or during memory usage, without the occurrence of endurance-limiting phenomena. As such, the accesses would be diverted from such blocks earlier than the remaining blocks, further insuring the integrity of the stored data.

In another embodiment of the invention, the functions of the directory processor 10 are combined in the counters. It is noted that the number stored in the counter is useless once a full count is obtained. The functions of the directory processor 10 can be combined in the counters by including, for example, three additional bits for each counter. These bits could indicate when the count is full, and by implication when other counter bits are to be used for other purposes. When the counter is full, as indicated by the stored

code in these three bits, then the other bits in the counter can serve as a pointer for indicating which substitutional block is to be addressed.

Claims

5

1. A nonvolatile storage means comprising:

a plurality of data storage arrays;

10

at least one counter storage array for storing values indicating the remaining useful life of each of said plurality of data storage arrays; and

15

means for controlling access to each of said plurality of data storage arrays as a function of the values stored in said counter array, said counter array having an appreciably longer useful life than any of said plurality of data storage arrays.

2. A nonvolatile memory system, comprising:

20

a plurality of nonvolatile memory blocks each having a plurality of memory cells, at least one of said plurality of nonvolatile memory blocks containing no data and not having been previously accessed;

a plurality of counters for counting a number of erase/write cycles that the plurality of nonvolatile memory blocks experience, respectively; and

25

means, coupled to said counters, for reallocating data contained in a first memory block to said at least one memory block when the counter for said first data block counts up to a predetermined value.

3. The system according to claim 2, wherein said predetermined value represents an endurance limit of nonvolatile memory cell blocks.

30

4. The system according to claim 2, wherein each of said plurality of counters comprises a non volatile binary counter and at least one nonvolatile linear counter.

35

5. The system according to claim 4, wherein each of said plurality of counters comprises two nonvolatile linear counters.

6. The system according to claim 5, wherein each of said plurality of counters further comprises an associated comparator means for comparing with each other the data contained in said two nonvolatile linear counters.

40

7. The system according to claim 6, further comprising means for incrementing one of said plurality of counters by one when its associated comparator means determines that the two nonvolatile linear counters are equal to each other, and for incrementing said one counter by a predetermined value which is substantially greater than one when said associated comparator means determines that said two nonvolatile linear counters are not equal to each other.

8. The system according to claim 2 or one of the claims 3 to 7, wherein each of said plurality of counters comprises a binary counter and a random number generator.

- 50 9. A method for counting a number of erase/write cycles of a nonvolatile memory system including a plurality of nonvolatile memory blocks, at least one of said plurality of nonvolatile memory blocks containing no data and not having been previously accessed, the method comprising:

55 counting a number of erase/write cycles that the plurality of nonvolatile memory blocks experience, respectively; and

reallocating data contained in a first memory block to said at least one memory block when the counted number of erase/write cycles of said first memory block is equal to a predetermined value.

10. The method according to claim 9, wherein said predetermined value represents an endurance limit of nonvolatile memory cells of said memory blocks.
- 5 11. The method according to claim 9 or 10, further comprising the steps of determining any errors in the nonvolatile memory system, and incrementing the counted number of erase/write cycles by one if no errors are determined.
- 10 12. The method according to claim 9, 10 or 11, further comprising the steps of determining any errors in the nonvolatile memory system, and incrementing the counted number of erase/write cycles by a number substantially greater than one if any errors are determined.
- 15 13. The method according to claim 9, 10, 11 or 12, further comprising the step of determining which nonvolatile memory blocks have a high endurance to erase/write cycles and which nonvolatile memory blocks have a low endurance to erase/write cycles.
14. The method according to claim 13, further comprising the step of preloading the counter associated with low endurance memory blocks with a count which is substantially greater than one.
- 20 15. The method according to claim 13 or 14, further comprising the step of designating at least one of high endurance memory blocks as said at least one memory block which contains no data.

25

30

35

40

45

50

55

FIG. 1

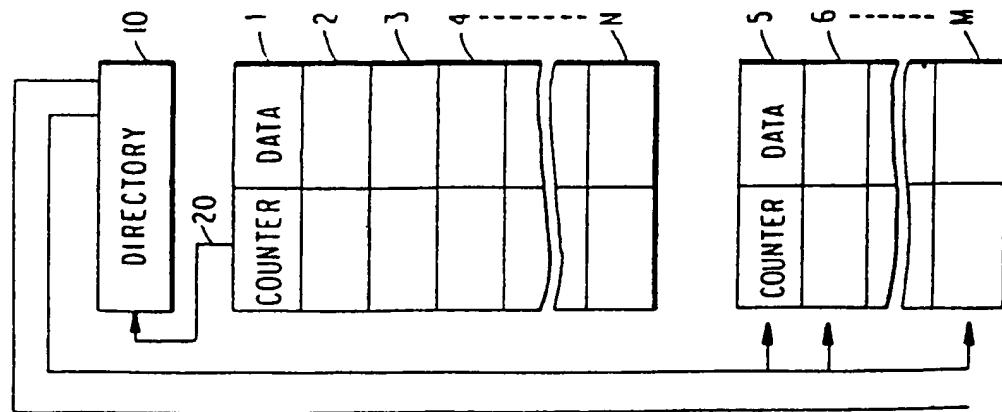


FIG. 2

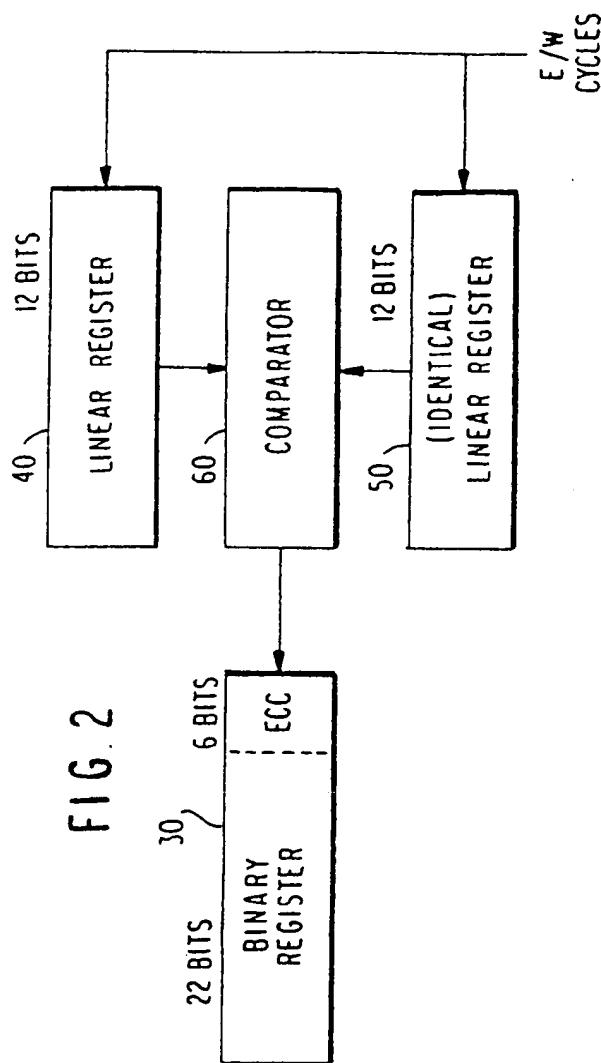
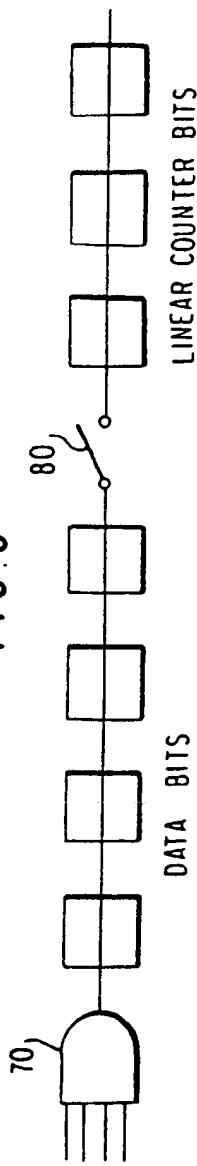


FIG. 3



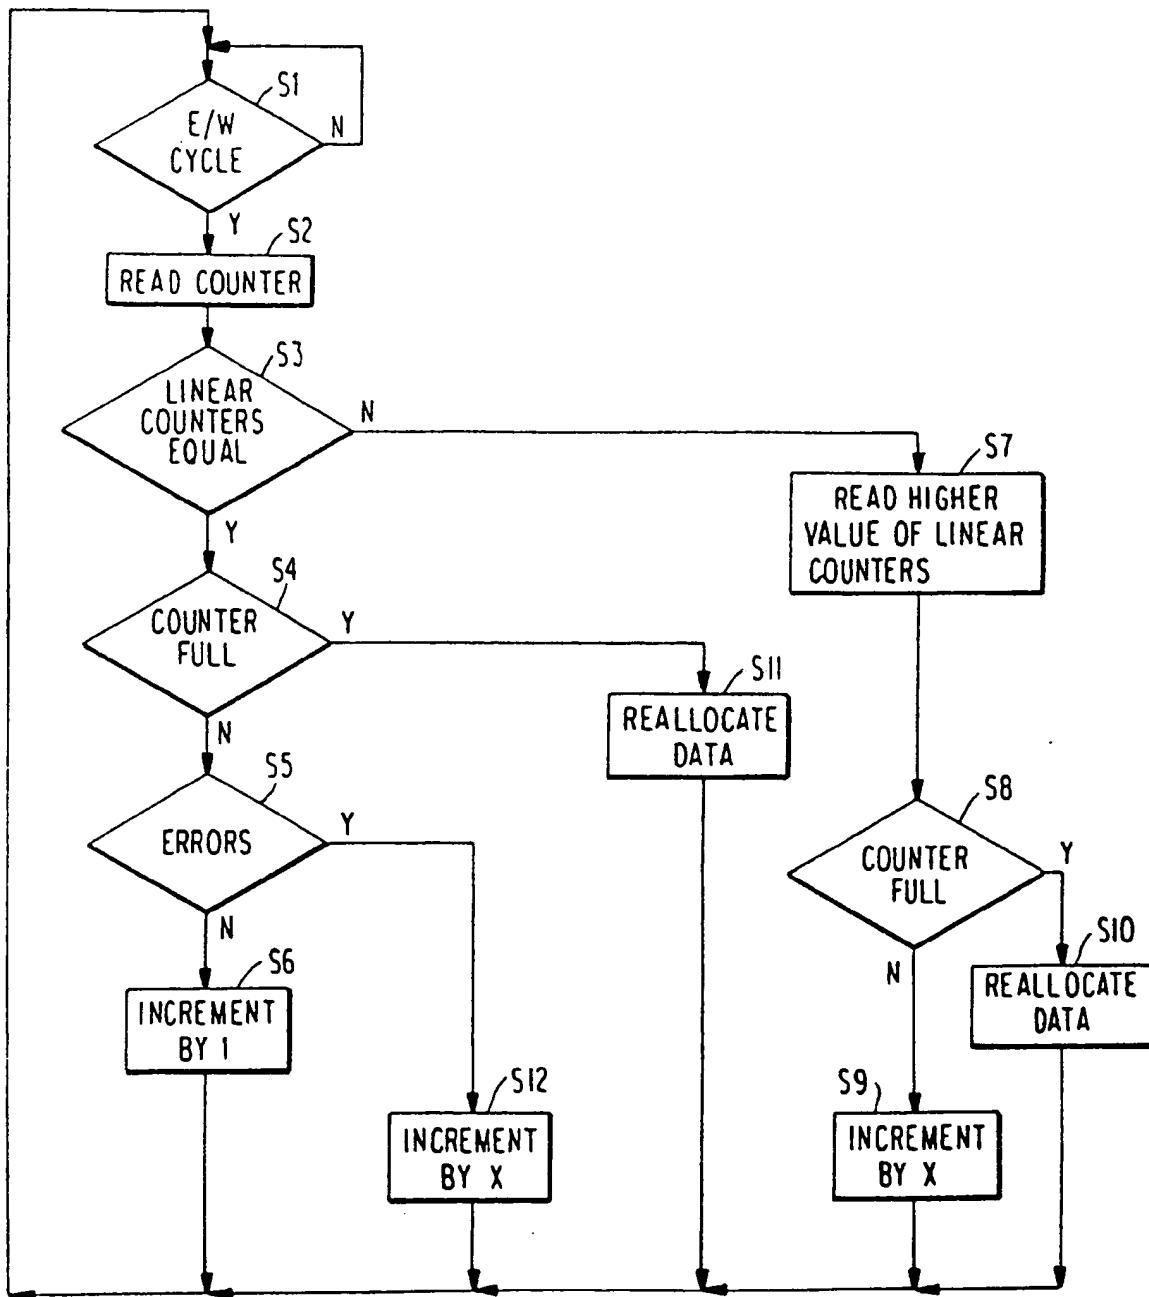


FIG. 4



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 91 11 9207

DOCUMENTS CONSIDERED TO BE RELEVANT			CLASSIFICATION OF THE APPLICATION (Int. Cl.5)						
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim							
Y	EP-A-0 398 654 (SEIKO EPSON CORPORATION) 22 November 1990 * column 2, line 57 - column 5, line 24 * ---	1-3, 9, 10	G01C22/02 G11C16/06 G06F11/20						
Y	US-A-4 612 640 (SEEQ TECHNOLOGY, INC.) 16 September 1986 * abstract * ---	1-3, 9, 10							
Y	PATENT ABSTRACTS OF JAPAN vol. 012, no. 173 (P-706) 24 May 1988 & JP-A-62 283 496 (CANON INC.) 9 December 1987 * abstract * ---	1-3, 9, 10							
A	EP-A-0 392 895 (SUNDISK CORPORATION) 17 October 1990 * column 9, line 35 - column 16, line 22 * ---	1							
A	DE-A-3 200 872 (SARTORIUS GMBH) 21 July 1983 ---								
			TECHNICAL FIELDS SEARCHED (Int. Cl.5)						
			G01C G11C G06F						
<p>The present search report has been drawn up for all claims</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 33%;">Place of search</td> <td style="width: 33%;">Date of completion of the search</td> <td style="width: 34%;">Examiner</td> </tr> <tr> <td>THE HAGUE</td> <td>26 MARCH 1992</td> <td>ABDALOM R.</td> </tr> </table>				Place of search	Date of completion of the search	Examiner	THE HAGUE	26 MARCH 1992	ABDALOM R.
Place of search	Date of completion of the search	Examiner							
THE HAGUE	26 MARCH 1992	ABDALOM R.							
CATEGORY OF CITED DOCUMENTS		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document							
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background : non-written disclosure P : intermediate document									